arXiv:1601.08051v1 [cs.DS] 29 Jan 2016

Minimal Suffix and Rotation of a Substring in Optimal Time*

Tomasz Kociumaka

Institute of Informatics, University of Warsaw, Poland
kociumaka@mimuw.edu.pl

Abstract

For a text given in advance, the substring minimal suffix queries ask to determine the lexicographically minimal non-empty suffix of a substring specified by the location of its occurrence in the text. We develop a data structure answering such queries optimally: in constant time after linear-time preprocessing. This improves upon the results of Babenko et al. (CPM 2014), whose trade-off solution is characterized by $\Theta(n \log n)$ product of these time complexities. Next, we extend our queries to support concatenations of $\mathcal{O}(1)$ substrings, for which the construction and query time is preserved. We apply these generalized queries to compute lexicographically minimal and maximal rotations of a given substring in constant time after linear-time preprocessing.

Our data structures mainly rely on properties of Lyndon words and Lyndon factorizations. We combine them with further algorithmic and combinatorial tools, such as fusion trees and the notion of order isomorphism of strings.

1 Introduction

Lyndon words, as well as the inherently linked concepts of the lexicographically minimal suffix and the lexicographically minimal rotation of a string, are one of the most successful concepts of combinatorics of words. Introduced by Lyndon [28] in the context of Lie algebras, they are widely used in algebra and combinatorics. They also have surprising algorithmic applications, including ones related to constant-space pattern matching [14], maximal repetitions [7], and the shortest common superstring problem [30].

The central combinatorial property of Lyndon words, proved by Chen et al. [9], states that every string can be uniquely decomposed into a non-increasing sequence of Lyndon words. Duval [16] devised a simple algorithm computing the Lyndon factorization in linear time and constant space. He also observed that the same algorithm can be used to determine the lexicographically minimal and maximal suffix, as well as the lexicographically minimal and maximal rotation of a given string.

The first two algorithms are actually on-line procedures: in linear time they allow computing the minimal and maximal suffix of every prefix of a given string. For rotations such a procedure was later introduced by Apostolico and Crochemore [3]. Both these solutions lead to the optimal, quadratic-time algorithms computing the minimal and maximal suffixes and rotations for all substring of a given string. Our main results are the data-structure versions of these problems: we preprocess a given text T to answer the following queries:

> Minimal Suffix Queries
> Given a substring $v = T[\ell..r]$ of T, report the lexicographically smallest non-empty suffix of v (represented by its length).

*This work is supported by Polish budget funds for science in 2013-2017 as a research project under the 'Diamond Grant' program.

MINIMAL ROTATION QUERIES
Given a substring $v = T[\ell..r]$ of T, report the lexicographically smallest rotation of v (represented by the number of positions to shift).

For both problems we obtain optimal solutions with linear construction time and constant query time. For MINIMAL SUFFIX QUERIES this improves upon the results of Babenko et al. [5], who developed a trade-off solution, which for a text of length n has $\Theta(n \log n)$ product of preprocessing and query time. We are not aware of any results for MINIMAL ROTATION QUERIES except for a data structure only testing cyclic equivalence of two subwords [26]. It allows constant-time queries after randomized preprocessing running in expected linear time.

An optimal solution for the MAXIMAL SUFFIX QUERIES was already obtained in [5], while the MAXIMAL ROTATION QUERIES are equivalent to MINIMAL ROTATION QUERIES subject to alphabet reversal. Hence, we do not focus on the maximization variants of our problems.

Using an auxiliary result devised to handle MINIMAL ROTATION QUERIES, we also develop a data structure answering in $\mathcal{O}(k^2)$ time the following generalized queries:

GENERALIZED MINIMAL SUFFIX QUERIES
Given a sequence of substrings $v_1, \ldots, v_k$ ($v_i = T[\ell_i..r_i]$), report the lexicographically smallest non-empty suffix of their concatenation $v_1 v_2 \ldots v_k$ (represented by its length).

All our algorithms are deterministic procedures for the standard word RAM model with machine words of size $W = \Omega(\log n)$ [19]. The alphabet is assumed to be $\Sigma = \{0, \ldots, \sigma - 1\}$ where $\sigma = n^{\mathcal{O}(1)}$, so that all letters of the input text T can be sorted in linear time.

Applications. The last factor of the Lyndon factorization of a string is its minimal suffix. As noted in [5], this can be used to reduce computing the factorization $v = v_1^{p_1} \cdots v_m^{p_m}$ of a substring $v = T[\ell..r]$ to $\mathcal{O}(m)$ MINIMAL SUFFIX QUERIES in T. Hence, our data structure determines the factorization in the optimal $\mathcal{O}(m)$ time. If v is a concatenation of k substrings, this increases to $\mathcal{O}(k^2 m)$ time (which we did not attempt to optimize in this paper).

The primary use of MINIMAL ROTATION QUERIES is *canonization* of substrings, i.e., classifying them according to cyclic equivalence (conjugacy); see [3]. As a proof-of-concept application of this natural tool, we propose counting distinct substring with a given exponent.

Related work. Our work falls in a class of substring queries: data structure problems solving basic stringology problems for substrings of a preprocessed text. This line of research, implicitly initiated by substring equality and longest common prefix queries (using suffix trees and suffix arrays; see [11]), now includes several problems related to compression [10, 24, 26, 6], pattern matching [26], and the range longest common prefix problem [1, 31, 2]. Closest to ours is a result by Babenko et al. [6], which after $\mathcal{O}(n\sqrt{\log n})$-expected-time preprocessing allows determining the k-th smallest suffix of a given substring, as well as finding the lexicographic rank of one substring among suffixes of another substring, both in logarithmic time

Outline of the paper. In Section 2 we recall standard definitions and two well-known data structures. Next, in Section 3, we study combinatorics of minimal suffixes, using in particular a notion of *significant suffixes*, introduced by I et al. [21, 22] to compute Lyndon factorizations of grammar-compressed strings. Section 4 is devoted to answering MINIMAL SUFFIX QUERIES. We use *fusion trees* by Pătraşcu and Thorup [32] to improve the query time from logarithmic to $\mathcal{O}(\log^* |v|)$, and then, by preprocessing shorts strings, we achieve constant query time. That final step uses a notion of *order-isomorphism* [27, 25] to reduce the number of precomputed values. In Section 5 we repeat the same steps for GENERALIZED MINIMAL SUFFIX QUERIES. We conclude with Section 6, where we briefly discuss the applications.

2 Preliminaries

We consider strings over an alphabet $\Sigma = \{0, \ldots, \sigma - 1\}$ with the natural order $\prec$. The empty string is denoted as ε. By Σ^* (Σ^+) we denote the set of all (resp. non-empty) finite strings over Σ. We also define Σ^∞ as the set of infinite strings over Σ. We extend the order $\prec$ on Σ in the standard way to the *lexicographic* order on $\Sigma^* \cup \Sigma^\infty$.

Let $w = w[1] \ldots w[n]$ be a string in Σ^*. We call n the *length* of w and denote it by $|w|$. For $1 \le i \le j \le n$, a string $u = w[i] \ldots w[j]$ is called a *substring* of w. By $w[i..j]$ we denote the occurrence of u at position i, called a *fragment* of w. A fragment of w other than the whole w is called a *proper* fragment of w. A fragment starting at position 1 is called a *prefix* of w and a fragment ending at position n is called a *suffix* of w. We use abbreviated notation $w[..j]$ and $w[i..]$ for a prefix $w[1..j]$ and a suffix $w[i..n]$ of w, respectively. A *border* of w is a substring of w which occurs both as a prefix and as a suffix of w. An integer p, $1 \le p \le |w|$, is a *period* of w if $w[i] = w[i+p]$ for $1 \le i \le n-p$. If w has period p, we also say that is has *exponent* $\frac{|w|}{p}$. Note that p is a period of w if and only if w has a border of length $|w| - p$.

We say that a string w' is a *rotation* (cyclic shift, conjugate) of a string w if there exists a decomposition $w = uv$ such that $w' = vu$. Here, w' is the left rotation of w by $|u|$ characters and the right rotation of w by $|v|$ characters.

Enhanced suffix array. The *suffix array* [29] of a text T of length n is a permutation SA of $\{1, \ldots, n\}$ defining the lexicographic order on suffixes $T[i..n]$: $SA[i] < SA[j]$ if and only if $T[i..n] \prec T[j..n]$. For a string T, both SA and its inverse permutation ISA take $\mathcal{O}(n)$ space and can be computed in $\mathcal{O}(n)$ time; see e.g. [11]. Typically, one also builds the LCP table and extends it with a data structure for range minimum queries [20, 8], so that the longest common prefix of any two suffixes of T can be determined efficiently.

Similarly to [5], we also construct these components for the reversed text T^R. Additionally, we preprocess the ISA table to answer range minimum and maximum queries. The resulting data structure, which we call the *enhanced suffix array* of T, lets us perform many queries.

Theorem 2.1 (Enhanced suffix array; see Fact 3 and Lemma 4 in [5])**.** *The enhanced suffix array of a text T of length n takes $\mathcal{O}(n)$ space, can be constructed in $\mathcal{O}(n)$ time, and allows answering the following queries in $\mathcal{O}(1)$ time given fragments x, y of T:*

(a) determine if $x \prec y$, $x = y$, or $x \succ y$,

(b) compute the the longest common prefix $\mathrm{lcp}(x, y)$ *and the longest common suffix* $\mathrm{lcs}(x, y)$*,*

(c) compute $\mathrm{lcp}(x^\infty, y)$ *and determine if $x^\infty \prec y$, $x^\infty = y$, or $x^\infty \succ y$.*

Moreover, given indices i, j, it can compute in $\mathcal{O}(1)$ time the minimal *and the* maximal *suffix among $\{T[k..n] : i \le k \le j\}$.*

Fusion trees. Consider a set $\mathcal{A}$ of W-bit integers (recall that W is the machine word size). *Rank* queries given a W-bit integer x return $\mathrm{rank}_\mathcal{A}(x)$ defined as $|\{y \in \mathcal{A} : y < x\}|$. Similarly, *select* queries given an integer r, $0 \le r < |\mathcal{A}|$, return $\mathrm{select}_\mathcal{A}(r)$, the r-th smallest element in $\mathcal{A}$, i.e., $x \in \mathcal{A}$ such that $\mathrm{rank}_\mathcal{A}(x) = r$. These queries can be used to determine the *predecessor* and the *successor* of a W-bit integer x, i.e., $\mathrm{pred}_\mathcal{A}(x) = \max\{y \in \mathcal{A} : y < x\}$ and $\mathrm{succ}_\mathcal{A}(x) = \min\{y \in \mathcal{A} : y \ge x\}$. We answer these queries with dynamic fusion trees by Pătraşcu and Thorup [32]. We only use these trees in a static setting, but the original static fusion trees by Fredman and Willard [17] do not have an efficient construction procedure.

Theorem 2.2 (Fusion trees [32, 17])**.** *There exists a data structure of size $\mathcal{O}(|\mathcal{A}|)$ which answers* $\mathrm{rank}_\mathcal{A}$*,* $\mathrm{select}_\mathcal{A}$*,* $\mathrm{pred}_\mathcal{A}$*, and* $\mathrm{succ}_\mathcal{A}$ *queries in $\mathcal{O}(1 + \log_W |\mathcal{A}|)$ time. Moreover, it can be constructed in $\mathcal{O}(|\mathcal{A}| + |\mathcal{A}| \log_W |\mathcal{A}|)$ time.*

3 Combinatorics of minimal suffixes and Lyndon words

For a non-empty string v the *minimal suffix* $\mathrm{MinSuf}(v)$ is the lexicographically smallest non-empty suffix s of v. Similarly, for an arbitrary string v the *maximal suffix* $\mathrm{MaxSuf}(v)$ is the lexicographically largest suffix s of v. We extend these notions as follows: for a pair of strings v, w we define $\mathrm{MinSuf}(v, w)$ and $\mathrm{MaxSuf}(v, w)$ as the lexicographically smallest (resp. largest) string sw such that s is a (possibly empty) suffix of v.

In order to relate minimal and maximal suffixes, we introduce the *reverse* order $\prec^R$ on Σ and extend it to the *reverse lexicographic* order, and an auxiliary symbol $\$ \notin \Sigma$. We extend the order $\prec$ on Σ so that $\mathtt{c} \prec \$$ (and thus $\$ \prec^R \mathtt{c}$) for every $\mathtt{c} \in \Sigma$. We define $\bar{\Sigma} = \Sigma \cup \{\$\}$, but unless otherwise stated, we still assume that the strings considered belong to Σ^*.

Observation 3.1. *If $u, v \in \Sigma^*$, then $u\$ \prec v$ if and only if $v \prec^R u$.*

We use MinSuf^R and MaxSuf^R to denote the minimal (resp. maximal) suffix with respect to $\prec^R$. The following observation relates the notions we introduced:

Observation 3.2. *(a)* $\mathrm{MaxSuf}(v, \varepsilon) = \mathrm{MaxSuf}(v)$ *for every* $v \in \bar{\Sigma}^*$,

(b) $\mathrm{MinSuf}(vw) = \min(\mathrm{MinSuf}(v, w), \mathrm{MinSuf}(w))$ *for every* $v \in \bar{\Sigma}^*$ *and* $w \in \bar{\Sigma}^+$,

(c) $\mathrm{MinSuf}(vc) = \mathrm{MinSuf}(v, c)$ *for every* $v \in \bar{\Sigma}^*$ *and* $c \in \bar{\Sigma}$,

(d) $\mathrm{MinSuf}(v, w\$) = \mathrm{MaxSuf}^R(v, w)\$$ *for every* $v, w \in \Sigma^*$,

(e) $\mathrm{MinSuf}(v\$) = \mathrm{MaxSuf}^R(v)\$$ *for every* $v \in \Sigma^*$.

A property seemingly similar to (e) is false for every $v \in \Sigma^+$: $\$ = \mathrm{MinSuf}^R(v\$) \neq \mathrm{MaxSuf}(v)\$$.

A notion deeply related to minimal and maximal suffixes is that of a Lyndon word [28, 9]. A string $w \in \Sigma^+$ is called a *Lyndon word* if $\mathrm{MinSuf}(w) = w$. Note that such w does not have proper borders, since a border would be a non-empty suffix smaller than w. A *Lyndon factorization* of a string $u \in \bar{\Sigma}^*$ is a representation $u = u_1^{p_1} \dots u_m^{p_m}$, where u_i are Lyndon words such that $u_1 \succ \dots \succ u_m$. Every non-empty word has a unique Lyndon factorization [9], which can be computed in linear time and constant space [16]. The following result provides a characterization of the Lyndon factorization of a concatenation of two strings:

Lemma 3.3 ([4, 15])**.** *Let $u = u_1^{p_1} \cdots u_m^{p_m}$ and $v = v_1^{q_1} \cdots v_\ell^{q_\ell}$ be Lyndon factorization. Then the Lyndon factorization of uv is $uv = u_1^{p_1} \cdots u_c^{p_c} z^k v_{d+1}^{q_{d+1}} \cdots v_\ell^{q_\ell}$ for integers c, d, k and a Lyndon word z such that $0 \le c < m$, $0 \le d \le \ell$, and $z^k = u_{c+1}^{p_{c+1}} \cdots u_m^{p_m} v_1^{q_1} \cdots v_d^{q_d}$.*

Next, we prove another simple yet useful property of Lyndon words:

Fact 3.4. *Let $v, w \in \Sigma^+$ be strings such that w is a Lyndon word. If $v \prec w$, then $v^\infty \prec w$.*

Proof. For a proof by contradiction suppose that $v \prec w \prec v^\infty$. Let $w = v^k s$, where v is not a prefix of s. Note that $k \ge 1$ as v must be a prefix of w. Because $w = v^k s \prec v^\infty$, we have $s \prec v^\infty$. On the other hand, w is a Lyndon word, so $v^k s = w \prec s$. Consequently, $v^k s \prec s \prec v^\infty$. Since $k \ge 1$, v must be a prefix of s, which contradicts the definition of k. □

3.1 Significant suffixes

Below we recall a notion of *significant suffixes*, introduced by I et al. [21, 22] in order to compute Lyndon factorizations of grammar-compressed strings. Then, we state combinatorial properties of significant suffixes; some of them are novel and some were proved in [22].

Definition 3.5 (see [21, 22])**.** A suffix s of a string $v \in \Sigma^*$ is a *significant suffix* of v if $sw = \mathrm{MinSuf}(v, w)$ for some $w \in \bar{\Sigma}^*$.

Let $v = v_1^{p_1} \dots v_m^{p_m}$ be the Lyndon factorization of a string $v \in \Sigma^+$. For $1 \le j \le m$ we denote $s_j = v_j^{p_j} \cdots v_m^{p_m}$; moreover, we assume $s_{m+1} = \varepsilon$. Let λ be the smallest index such that s_{i+1} is a prefix of v_i for $\lambda \le i \le m$. Observe that $s_\lambda \succ \ldots \succ s_m \succ s_{m+1} = \varepsilon$, since v_i is a prefix of s_i. We define y_i so that $v_i = s_{i+1}y_i$, and we set $x_i = y_i s_{i+1}$. Note that $s_i = v_i^{p_i} s_{i+1} = (s_{i+1}y_i)^{p_i} s_{i+1} = s_{i+1}(y_i s_{i+1})^{p_i} = s_{i+1}x_i^{p_i}$. We also denote $\Lambda(w) = \{s_\lambda, \dots, s_m, s_{m+1}\}$, $X(w) = \{x_\lambda^\infty, \dots, x_m^\infty\}$, and $X'(w) = \{x_\lambda^{p_\lambda}, \dots, x_m^{p_m}\}$. The observation below lists several immediate properties of the introduced strings:

Observation 3.6. *For each i, $\lambda \le i \le m$: (a) $x_i^\infty \succ x_i^{p_i} \succeq x_i \succeq y_i$, (b) $x_i^{p_i}$ is a suffix of v of length $|s_i| - |s_{i+1}|$, and (c) $|s_i| > 2|s_{i+1}|$. In particular, $|\Lambda(v)| = \mathcal{O}(\log |v|)$.*

The following lemma shows that $\Lambda(v)$ is equal to the set of significant suffixes of v. (Significant suffixes are actually defined in [22] as $\Lambda(v)$ and only later proved to satisfy our Definition 3.5.) In fact, the lemma is much deeper; in particular, the formula for $\mathrm{MaxSuf}(v, w)$ is one of the key ingredients of our efficient algorithms answering Minimal Suffix Queries.

Lemma 3.7 (I et al. [22], Lemmas 12–14)**.** *For a string $v \in \Sigma^+$ let s_i, λ, x_i, and y_i, be defined as above. Then $x_\lambda^\infty \succ x_\lambda^{p_\lambda} \succeq y_\lambda \succ x_{\lambda+1}^\infty \succ x_{\lambda+1}^{p_{\lambda+1}} \succeq y_{\lambda+1} \succ \ldots \succ x_m^\infty \succ x_m^{p_m} \succeq y_m$. Moreover, for every string $w \in \bar{\Sigma}^*$ we have*

$$\mathrm{MinSuf}(v, w) = \begin{cases} s_\lambda w & \text{if } w \succ x_\lambda^\infty, \\ s_i w & \text{if } x_{i-1}^\infty \succ w \succ x_i^\infty \text{ for } \lambda < i \le m, \\ s_{m+1} w & \text{if } x_m^\infty \succ w. \end{cases}$$

In other words, $\mathrm{MinSuf}(v, w) = s_{m+1-r} w$ *where* $r = \mathrm{rank}_{X(v)}(w)$.

We apply Lemma 3.7 to deduce several properties of the set $\Lambda(v)$ of significant suffixes.

Corollary 3.8. *For every string $v \in \Sigma^+$:*

(a) the largest suffix in $\Lambda(v)$ is $\mathrm{MaxSuf}^R(v)$ and $\Lambda(v) = \Lambda(\mathrm{MaxSuf}^R(v))$,

(b) if s is a suffix of v such that $|v| \le 2|s|$, then $\Lambda(v) \subseteq \Lambda(s) \cup \{\mathrm{MaxSuf}^R(v)\}$.

Proof. To prove (a), observe that $x_\lambda \in \Sigma^+$, so $x_\lambda^\infty \prec \$$. Consequently, Lemma 3.7 states that $s_\lambda\$ = \mathrm{MinSuf}(v, \$)$. However, we have $\mathrm{MinSuf}(v, \$) = \mathrm{MaxSuf}^R(v)\$$ by Observation 3.2(e), and thus $s_\lambda = \mathrm{MaxSuf}^R(v)$. Uniqueness of the Lyndon factorization implies that $u_\lambda^{p_\lambda} \cdots u_m^{p_m}$ is the Lyndon factorization of s_λ, and hence by definition of $\Lambda(\cdot)$ we have $\Lambda(v) = \Lambda(s_\lambda)$.

For a proof of (b), we shall show that for $i \ge \lambda + 1$ the string s_i is a significant suffix of s. Note that, by Observation 3.6, s_i is a suffix of s, since $2|s_i| < |s_{i-1}| \le |s_\lambda| \le |v| \le 2|s|$. The suffix $s_{m+1} = \varepsilon$ is clearly a significant suffix of s, so we assume $\lambda < i \le m$. By Lemma 3.7, one can choose $w \in \bar{\Sigma}^*$ (setting $x_{i-1}^\infty \succ w \succ x_i^\infty$) so that $s_i w = \mathrm{MinSuf}(v, w)$. However, this also implies $s_i w = \mathrm{MinSuf}(s, w)$ because all suffixes of s are suffixes of v. Consequently, s_i is a significant suffix of s, as claimed. □

Below we provide a precise characterization of $\Lambda(uv)$ for $|u| \le |v|$ in terms of $\Lambda(v)$ and $\mathrm{MaxSuf}^R(u, v)$. This is another key ingredient of our data structure, in particular letting us efficiently compute significant suffixes of a given fragment of T.

Lemma 3.9. *Let $u, v \in \Sigma^+$ be strings such that $|u| \le |v|$. Also, let $\Lambda(v) = \{s_\lambda, \dots, s_{m+1}\}$, $s' = \mathrm{MaxSuf}^R(u, v)$, and let s_i be the longest suffix in $\Lambda(v)$ which is a prefix of s'. Then*

$$\Lambda(uv) = \begin{cases} \{s_\lambda, \dots, s_{m+1}\} & \text{if } s' \preceq^R s_\lambda \text{ (i.e., if } s_\lambda \preceq s' \text{ and } i \ne \lambda), \\ \{s', s_{i+1}, \dots, s_{m+1}\} & \text{if } s' \succ^R s_\lambda,\ i \le m, \text{ and } |s_i| - |s_{i+1}| \text{ is a period of } s', \\ \{s', s_i, s_{i+1}, \dots, s_{m+1}\} & \text{otherwise.} \end{cases}$$

Consequently, for every $w \in \bar{\Sigma}^*$, we have $\mathrm{MinSuf}(uv, w) \in \{\mathrm{MaxSuf}^R(u, v)w, \mathrm{MinSuf}(v, w)\}$.

Proof. Observation 3.2 yields $\mathrm{MaxSuf}^R(uv) \in \{\mathrm{MaxSuf}^R(u,v), \mathrm{MaxSuf}^R(v)\}$. By Corollary 3.8(a) this is equivalent to $\mathrm{MaxSuf}^R(uv) \in \{s', s_\lambda\}$. Consequently, if $s' \preceq^R s_\lambda$, then $\mathrm{MaxSuf}^R(uv) = s_\lambda$ and Corollary 3.8(a) implies $\Lambda(uv) = \Lambda(s_\lambda) = \Lambda(v)$, as claimed.

Thus, we may assume that $s' \succ^R s_\lambda$, and in particular that $s' = \mathrm{MaxSuf}^R(uv)$. Let $s_j \in \Lambda(w)$ be the longest suffix in $\Lambda(uv) \cap \Lambda(v)$ ($\lambda \leq j \leq m+1$). By Corollary 3.8(b), $\Lambda(uv) \subseteq \{s'\} \cup \{s_j, s_{j+1}, \ldots, s_{m+1}\}$. Lemma 3.3 and the definition the $\Lambda(\cdot)$ set in terms of the Lyndon factorization yield that the inclusion above is actually an equality. Moreover, the definition also implies that s_j is a prefix of s', and thus $j \geq i$. If $i = m+1$, this already proves our statement, so in the remainder of the proof we assume $i \leq m$.

First, let us suppose that $j \geq i+1$. We shall prove that $j = i+1$ and $|s_i| - |s_{i+1}|$ is a period of s'. Let u' be a string such that $s' = u's_j$. Note that $v_i^{p_i} \ldots v_{j-1}^{p_{j-1}}$ is a border of u' (as s_i is a border of s'), so v_{j-1} is also a border of u' (because v_{j-1} is a prefix of s_{j-1}, which is a prefix of v_i). Moreover, by definition of the $\Lambda(uv)$ set, u' must be a power of a Lyndon word. Lyndon words do not proper borders, so any border of u' must be a power of the same Lyndon word. Thus, u' is a power of v_{j-1}. As v_i is a Lyndon word and a prefix of u', this means that $|v_i| \leq |v_{j-1}|$. Consequently, $i = j+1$ since $|v_i| > |v_{i+1}| > \ldots > |v_m|$. What is more, as s_{i+1} is a prefix of v_i, we conclude that $|v_i|$ is a period of $s' = u's_{i+1}$. Therefore, $|s_i| - |s_{i+1}| = p_i|v_i|$ is also a period of s'.

It remains to prove that $j = i$ implies that $|s_i| - |s_{i+1}|$ is not a period of s'. For a proof by contradiction suppose that both $s_i \in \Lambda(uv)$ and $|s_i| - |s_{i+1}| = p_i|v_i|$ is a period of s'. Let us define u' so that $s' = u's_i = u'v_i^{p_i}s_{i+1}$. As $p_i|v_i|$ is a period of s' and $v_i^{p_i}$ contained in s', we conclude that s' is a substring of $(v_i^{p_i})^\infty = v_i^\infty$, and consequently $|v_i|$ is also a period of s' and hence a period of u' as well. However, by definition of the $\Lambda(\cdot)$ set, u' is a power of a Lyndon word whose length exceeds $|s_i|$ and thus also $|v_i|$. This Lyndon word cannot have a proper border, and such a border is induced by period $|v_i|$, a contradiction.

Finally, observe that the second claim easily follows from $\Lambda(uv) \subseteq \Lambda(v) \cup \{s'\}$. □

We conclude with two combinatorial lemmas, useful to in determining $\mathrm{MaxSuf}^R(u,v)$ for $|u| \leq |v|$. The first of them is also applied later in Section 5.

Lemma 3.10. *Let $v \in \Sigma^+$ and $w, w' \in \bar{\Sigma}^+$ be strings such that $w \prec w'$ and the longest common prefix of w and w' is not a proper substring of v. Also, let $\Lambda(v) = \{s_\lambda, \ldots, s_{m-1}\}$. If $\mathrm{MinSuf}(v,w) = s_iw$, then $\mathrm{MinSuf}(v,w') \in \{s_{i-1}w', s_iw'\}$.*

Proof. Due to the characterization in Lemma 3.7, we may equivalently prove that $\mathrm{rank}_{X(v)}(w')$ is $\mathrm{rank}_{X(v)}(w)$ or $\mathrm{rank}_{X(v)}(w)+1$. Clearly, $\mathrm{rank}_{X(v)}(w) \leq \mathrm{rank}_{X(v)}(w')$, so it suffices to show that $\mathrm{rank}_{X(v)}(w') \leq \mathrm{rank}_{X(v)}(w)+1$. This is clear if $|X(v)| = 1$, so we assume $|X(v)| > 1$.

This assumption in particular yields that $X'(v)$ consists of proper substrings of v, and thus $\mathrm{rank}_{X'(v)}(w) = \mathrm{rank}_{X(v)}(w')$ by the condition on the longest common prefix of w and w'. However, the inequality in Lemma 3.7 implies

$$\mathrm{rank}_{X(v)}(w') \leq \mathrm{rank}_{X'(v)}(w') = \mathrm{rank}_{X'(v)}(w) \leq \mathrm{rank}_{X(v)}(w)+1.$$

This concludes the proof. □

Lemma 3.11. *Let $v \in \Sigma^+$, $v = v_1^{p_1} \cdots v_m^{p_m}$ be the Lyndon factorization of v, and let $\Lambda(v) = \{s_\lambda, \ldots, s_{m+1}\}$. If for some $w \in \bar{\Sigma}^*$ and $\lambda < i \leq m+1$ we have $\mathrm{MinSuf}(v,w) = s_iw$, then $v_{i-1}s_iw \preceq sw$ for every non-empty suffix s of v satisfying $|s| > |s_i|$.*

Proof. Let s' be a string such that $s = s's_i$. First, suppose that $|s| < |v_{i-1}s_i|$. In this case s' is a proper suffix of a Lyndon word v_{i-1}, and thus $s' \succ v_i$ and, moreover, $sw \succ s' \succ v_{i-1}s_iw$. Thus, we may assume that $|s| > |v_{i-1}s_i|$.

Let $w' = v_{i-1}s_iw$ and let v' be a string such that $v = v'v_{i-1}s_i$. Observe that it suffices to prove that $\mathrm{MinSuf}(v', w') = w'$, which implies that $sw \preceq w'$ for $|s| > |v_{i-1}s_i|$. If $v' = \varepsilon$ there is nothing to prove, so we shall assume $|v'| > 0$. Note that we have the Lyndon factorization $v' = v_1^{p_1} \cdots v_{i-1}^{p_{i-1}-1}$ with $i > 1$ or

$p_{i-1} > 1$. By Lemma 3.7, $\mathrm{MinSuf}(v, w) = s_i w$ implies $w \prec x_{i-1}^{\infty}$ and $\mathrm{MinSuf}(v', w') = w'$ is equivalent to $w' \prec v_{i-1}^{\infty}$ (if $p_{i-1} > 1$) or $w' \prec v_{i-2}^{\infty}$ (if $p_{i-1} = 1$). We have

$$w' = v_{i-1}s_i w \prec v_{i-1}s_i x_{i-1}^{\infty} = v_{i-1}s_i(y_{i-1}s_i)^{\infty} = v_{i-1}(s_i y_{i-1})^{\infty} = v_{i-1}v_{i-1}^{\infty} = v_{i-1}^{\infty}$$

as claimed. If $p_{i-1} > 1$, this already concludes the proof, and thus we may assume that $p_{i-1} = 1$. By definition of the Lyndon factorization we have $v_{i-2} \succ v_{i-1}$, and by Fact 3.4 this implies $v_{i-2} \succ v_{i-1}^{\infty}$. Hence, $w' \prec v_{i-1}^{\infty} \prec v_{i-2} \prec v_{i-2}^{\infty}$, which concludes the proof. □

4 Answering Minimal Suffix Queries

In this section we present our data structure for Minimal Suffix Queries. We proceed in three steps improving the query time from $\mathcal{O}(\log |v|)$ via $\mathcal{O}(\log^* |v|)$ to $\mathcal{O}(1)$. The first solution is an immediate application of Observation 3.2(c) and the notion of significant suffixes. Efficient computation of these suffixes, also used in the construction of further versions of our data structure, is based on Lemma 3.9, which yields a recursive procedure. The only "new" suffix needed at each step is determined using the following result, which can be seen as a cleaner formulation of Lemma 14 in [5].

Lemma 4.1. *Let $u = T[\ell..r]$ and $v = T[r+1..r']$ be fragments of T such that $|u| \le |v|$. Using the enhanced suffix array of T we can compute $\mathrm{MaxSuf}^R(u, v)$ in $\mathcal{O}(1)$ time.*

Proof. Let $sv = \mathrm{MaxSuf}^R(u, v)$ and note that, by Observation 3.2(d), $sv\$ = \mathrm{MinSuf}(u, v\$)$. Let us focus on determining the latter value. The enhanced suffix array lets us compute a index k, $\ell \le k \le r$, which minimizes $T[k..]$. Equivalently, we have $T[k..] = \mathrm{MinSuf}(u, T[r+1..])$. Consequently, $T[k..r] = s_i \in \Lambda(u)$ for some $\lambda \le i \le m+1$. Since $|u| \le |v|$, v is not a proper substring of u, and by Lemma 3.10, we have $s \in \{s_{i-1}, s_i\}$ (if $i = \lambda$, then $s = s_i$).

Thus, we shall generate a suffix of s_{i-1} equal to s_{i-1} if $i > \lambda$, and return the better of the two candidates for $\mathrm{MinSuf}(u, v\$)$. If $k = \ell$, we must have $i = \lambda$ and there is nothing to do. Hence, let us assume $k > \ell$. By Lemma 3.11, if we compute an index k', $\ell \le k' < k$, which minimizes $T[k'..]$, we shall have $T[k'..k-1] = u_{i-1}$ provided that $i > \lambda$. Now, p_{i-1} can be generated as the largest integer such that $u_{i-1}^{p_{i-1}}$ is a suffix of $T[\ell..k-1]$, and we have $|s_{i-1}| = |s_i| + p_{i-1}|u_{i-1}|$, which lets us determine s_{i-1}. □

Lemma 4.2. *Given a fragment v of T, we can compute $\Lambda(v)$ in $\mathcal{O}(\log |v|)$ time using the enhanced suffix array of T*

Proof. If $|v| = 1$, we return $\Lambda(v) = \{v, \varepsilon\}$. Otherwise, we decompose $v = uv'$ so that $|v'| = \lceil \frac{1}{2}|v| \rceil$. We recursively generate $\Lambda(v')$ and use Lemma 4.1 to compute $s = \mathrm{MaxSuf}^R(u, v')$. Then, we apply the characterization of Lemma 3.9 to determine $\Lambda(v) = \Lambda(uv')$, using the enhanced suffix array (Theorem 2.1) to lexicographically compare fragments of T.

We store the lengths of the significant suffixes in an ordered list. This way we can implement a single phase (excluding the recursive calls) in time proportional to $\mathcal{O}(1)$ plus the number of suffixes removed from $\Lambda(v')$ to obtain $\Lambda(v)$. Since this is amortized constant time, the total running time becomes $\mathcal{O}(\log |v|)$ as announced. □

Corollary 4.3. Minimal Suffix Queries *queries can be answered in $\mathcal{O}(\log |v|)$ time using the enhanced suffix array of T.*

Proof. Recall that Observation 3.2(c) yields $\mathrm{MinSuf}(v) = \mathrm{MinSuf}(v[1..m-1], v[m])$ where $m = |v|$. Consequently, $\mathrm{MinSuf}(v) = sv[m]$ for some $s \in \Lambda(v[1..m-1])$. We apply Lemma 4.2 to compute $\Lambda(v[1..m-1])$ and determine the answer among $\mathcal{O}(\log |v|)$ candidates using lexicographic comparison of fragments, provided by the enhanced suffix array (Theorem 2.1). □

4.1 $\mathcal{O}(\log^* |v|)$-time Minimal Suffix Queries

An alternative $\mathcal{O}(\log |v|)$-time algorithm could be developed based just on the second part of Lemma 3.9: decompose $v = uv'$ so that $|v'| > |u|$ and return $\min(\mathrm{MaxSuf}^R(u, v'), \mathrm{MinSuf}(v'))$. The result is $\mathrm{MinSuf}(v)$ due to Lemma 3.9 and Observation 3.2(c). Here, the first candidate $\mathrm{MaxSuf}^R(u, v')$ is determined via Lemma 4.1, while the second one using a recursive call. A way to improve query time to $\mathcal{O}(1)$ at the price of $\mathcal{O}(n \log n)$-time preprocessing is to precompute the answers for *basic* fragments, i.e., fragments whose length is a power of two. Then, in order to determine $\mathrm{MinSuf}(v)$, we perform just a single step of the aforementioned procedure, making sure that v' is a basic fragment. Both these ideas are actually present in [5], along with a smooth trade-off between their preprocessing and query times.

Our $\mathcal{O}(\log^* |v|)$-time query algorithm combines recursion with preprocessing for certain *distinguished* fragments. More precisely, we say that $v = T[\ell..r]$ is distinguished if both $|v| = 2^q$ and $f(2^q) \mid r$ for some positive integer q, where $f(x) = 2^{\lfloor \log \log x \rfloor^2}$. Note that the number of distinguished fragments of length 2^q is at most $\frac{n}{2^{\lfloor \log q \rfloor^2}} = \mathcal{O}(\frac{n}{q^{\omega(1)}})$.

The query algorithm is based on the following decomposition ($x > f(x)$ for $x > 2^{16}$):

Fact 4.4. *Given a fragment $v = T[\ell..r]$ such that $|v| > f(|v|)$, we can in constant time decompose $v = uv'v''$ such that $1 \le |v''| \le f(|v|)$, v' is distinguished, and $|u| \le |v'|$.*

Proof. Let $q = \lfloor \log |v| \rfloor$ and $q' = \lfloor \log q \rfloor^2$. We determine r' as the largest integer strictly smaller than r divisible by $2^{q'} = f(|v|)$. By the assumption that $|v| > 2^{q'}$, we conclude that $r' \ge r - 2^{q'} \ge \ell$. We define $v'' = T[r' + 1..r]$ and partition $T[\ell..r'] = uv'$ so that $|v'|$ is the largest possible power of two. This guarantees $|u| \le |v'|$. Moreover, $|v'| \le |v|$ assures that $f(|v'|) \mid f(|v|)$, so $f(|v'|) \mid r'$, and therefore v' is indeed distinguished. □

Observation 3.2(b) implies that $\mathrm{MinSuf}(v) \in \{\mathrm{MinSuf}(uv', v''), \mathrm{MinSuf}(v'')\}$. Lemma 3.9 further yields $\mathrm{MinSuf}(v) \in \{\mathrm{MaxSuf}^R(u, v')v'', \mathrm{MinSuf}(v', v''), \mathrm{MinSuf}(v'')\}$. In other words, it leaves us with three candidates for $\mathrm{MinSuf}(v)$. Our query algorithm obtains $\mathrm{MaxSuf}^R(u, v')$ using Lemma 4.1, computes $\mathrm{MinSuf}(v'')$ recursively, and determines $\mathrm{MinSuf}(v', v'')$ through the characterization of Lemma 3.7. The latter step is performed using the following component based on a fusion tree, which we build for all distinguished fragments.

Lemma 4.5. *Let $v = T[\ell..r]$ be a fragment of T. There exists a data structure of size $\mathcal{O}(\log |v|)$ which answers the following queries in $\mathcal{O}(1)$ time: given a position $r' > r$ compute $\mathrm{MinSuf}(v, T[r + 1..r'])$. Moreover, this data structure can be constructed in $\mathcal{O}(\log |v|)$ time using the enhanced suffix array of T.*

Proof. By Lemma 3.7, we have $\mathrm{MinSuf}(v, w) = s_{m+1-\mathrm{rank}_{X(v)}(w)}w$, so in order to determine $\mathrm{MinSuf}(v, T[r + 1..r'])$, it suffices to store $\Lambda(v)$ and efficiently compute $\mathrm{rank}_{X(v)}(w)$ given $w = T[r + 1..r']$. We shall reduce these rank queries to rank queries in an integer set $R(v)$.

Claim. Denote $X(v) = \{x_\lambda^\infty, \ldots, x_m^\infty\}$ and let

$$R(v) = \{r + \mathrm{lcp}(T[r + 1..], x_j^\infty) : x_j^\infty \in X(w) \wedge x_j^\infty \prec T[r + 1..]\}.$$

For every index r', $r < r' \le n$, we have $\mathrm{rank}_{X(v)}(T[r + 1..r']) = \mathrm{rank}_{R(v)}(r')$.

Proof. We shall prove that for each j, $\lambda \le j \le m$, we have

$$x_j^\infty \prec T[r + 1..r'] \iff \big(r + \mathrm{lcp}(T[r + 1..], x_j^\infty) < r' \ \wedge \ x_j^\infty \prec T[r + 1..]\big).$$

First, if $x_j^\infty \succ T[r+1..]$, then clearly $x_j^\infty \succ T[r+1..r']$ and both sides of the equivalence are false. Therefore, we may assume $x_j^\infty \prec T[r+1..]$. Observe that in this case $d := \mathrm{lcp}(T[r+1..], x_j^\infty)$ is strictly less than $n-r$, and $T[r + 1..r + d] \prec x_j^\infty \prec T[r + 1..r + d + 1]$. Hence, $x_j^\infty \prec T[r + 1..r']$ if and only if $r + d < r'$, as claimed. □

We apply Theorem 2.2 to build a fusion tree for $R(v)$, so that the ranks are can be obtained in $\mathcal{O}(1 + \frac{\log|R(v)|}{\log W})$ time, which is $\mathcal{O}(1 + \frac{\log\log|v|}{\log\log n}) = \mathcal{O}(1)$ by Observation 3.6.

The construction algorithm uses Lemma 4.2 to compute $\Lambda(v) = \{s_\lambda, \ldots, s_{m+1}\}$. Next, for each j, $\lambda \le j \le m$, we need to determine $\mathrm{lcp}(T[r+1..], x_j^\infty)$. This is the same as $\mathrm{lcp}(T[r+1..], (x_j^{p_j})^\infty)$ and, by Observation 3.6, $x_j^{p_j}$ can be retrieved as the suffix of v of length $|s_i| - |s_{i+1}|$. Hence, the enhanced suffix array can be used to compute these longest common prefixes and therefore to construct $R(v)$ in $\mathcal{O}(|\Lambda(v)|) = \mathcal{O}(\log|v|)$ time. □

With this central component we are ready to give a full description of our data structure.

Theorem 4.6. *For every text T of length n there exists a data structure of size $\mathcal{O}(n)$ which answers* Minimal Suffix Queries *in $\mathcal{O}(\log^*|v|)$ time and can be constructed in $\mathcal{O}(n)$ time.*

Proof. Our data structure consists of the enhanced suffix array (Theorem 2.1) and the components of Lemma 4.5 for all distinguished fragments of T. Each such fragment of length 2^q contributes $\mathcal{O}(q)$ to the space consumption and to the construction time, which in total over all lengths sums up to $\mathcal{O}(\sum_q \frac{nq}{q^{\omega(1)}}) = \mathcal{O}(\sum_q \frac{n}{q^{\omega(1)}}) = \mathcal{O}(n)$.

Let us proceed to the query algorithm. Assume we are to compute the minimal suffix of a fragment v. If $|v| \le f(|v|)$ (i.e., if $|v| \le 2^{16}$), we use the logarithmic-time query algorithm given in Corollary 4.3. If $|v| > 2^q$, we apply Fact 4.4 to determine a decomposition $v = uv'v''$, which gives us three candidates for $\mathrm{MinSuf}(v)$. As already described, $\mathrm{MinSuf}(v'')$ is computed recursively, $\mathrm{MinSuf}(v', v'')$ using Lemma 4.5, and $\mathrm{MaxSuf}^R(u, v')v''$ using Lemma 4.1. The latter two both support constant-time queries, so the overall time complexity is proportional to the depth of the recursion. We have $|v''| \le f(|v|) < |v|$, so it terminates. Moreover,

$$f(f(x)) = 2^{\lfloor\log(\log f(x))\rfloor^2} \le 2^{(\log(\log\log x)^2)^2} = 2^{4(\log\log\log x)^2} = 2^{o(\log\log x)} = o(\log x).$$

Thus, $f(f(x)) \le \log x$ unless $x = \mathcal{O}(1)$. Consequently, unless $|v| = \mathcal{O}(1)$, when the algorithm clearly needs constant time, the length of the queried fragment is in two steps reduced from $|v|$ to at most $\log|v|$. This concludes the proof that the query time is $\mathcal{O}(\log^*|v|)$. □

4.2 $\mathcal{O}(1)$-time Minimal Suffix Queries

The $\mathcal{O}(\log^*|v|)$ time complexity of the query algorithm of Theorem 4.6 is only due to the recursion, which in a single step reduces the length of the queried fragment from $|v|$ to $f(|v|)$ where $f(x) = 2^{\lfloor\log\log x\rfloor^2}$. Since $f(f(x)) = 2^{o(\log\log x)}$, after just two steps the fragment length does not exceed $f(f(n)) = o(\frac{\log n}{\log\log n})$. In this section we show that the minimal suffixes of such short fragments can precomputed in a certain sense, and thus after reaching $\tau = f(f(n))$ we do not need to perform further recursive calls.

For constant alphabets, we could actually store all the answers for all $\mathcal{O}(\sigma^\tau) = n^{o(1)}$ strings of length up to τ. Nevertheless, in general all letters of T, and consequently all fragments of T, could even be distinct. However, the answers to Minimal Suffix Queries actually depend only on the relative order between letters, which is captured by order-isomorphism.

Two strings x and y are called *order-isomorphic* [27, 25], denoted as $x \approx y$, if $|x| = |y|$ and for every two positions i, j ($1 \le i, j \le |x|$) we have $x[i] \prec x[j] \iff y[i] \prec y[j]$. Note that the equivalence extends to arbitrary corresponding fragments of x and y, i.e., $x[i..j] \prec x[i'..j'] \iff y[i..j] \prec y[i'..j']$. Consequently, order-isomorphic strings cannot be distinguished using Minimal Suffix Queries or Generalized Minimal Suffix Queries.

Moreover, note that every string of length m is order-isomorphic to a string over an alphabet $\{1, \ldots, m\}$. Consequently, order-isomorphism partitions strings of length up to m into $\mathcal{O}(m^m)$ equivalence classes. The following fact lets us compute canonical representations of strings whose length is bounded by $m = W^{\mathcal{O}(1)}$.

Fact 4.7. *For every fixed integer $m = W^{\mathcal{O}(1)}$, there exists a function* oid *mapping each string w of length up to m to a non-negative integer* $\mathrm{oid}(w)$ *with $\mathcal{O}(m\log m)$ bits, so that $w \approx w' \iff \mathrm{oid}(w) = \mathrm{oid}(w')$. Moreover, the function can be evaluated in $\mathcal{O}(m)$ time.*

Proof. To compute $\mathrm{oid}(w)$, we first build a fusion tree storing all (distinct) letters which occur in w. Next, we replace each character of w with its rank among these letters. We allocate $\lceil \log m \rceil$ bits per character and prepend such a representation with $\lceil \log m \rceil$ bits encoding $|w|$. This way $\mathrm{oid}(w)$ is a sequence of $(|w| + 1)\lceil \log m \rceil = \mathcal{O}(m \log m)$ bits. Using Theorem 2.2 to build the fusion tree, we obtain an $\mathcal{O}(m)$-time evaluation algorithm. □

To answer queries for short fragments of T, we define overlapping *blocks* of length $m = 2\tau$: for $0 \leq i \leq \frac{n}{\tau}$ we create a block $T_i = T[1 + i\tau .. \min(n, (i+2)\tau)]$. For each block we apply Fact 4.7 to compute the identifier $\mathrm{oid}(T_i)$. The total length of the blocks is bounded $2n$, so this takes $\mathcal{O}(n)$ time. The identifiers use $\mathcal{O}(\frac{n}{\tau}\tau \log \tau) = O(n \log \tau)$ bits of space.

Moreover, for each distinct identifier $\mathrm{oid}(T_i)$, we store the answers to all the Minimal Suffix Queries queries in T_i. This takes $\mathcal{O}(\log m)$ bits per answer, and $\mathcal{O}(2^{\mathcal{O}(m \log m)} m^2 \log m) = 2^{\mathcal{O}(\tau \log \tau)}$ in total. Since $\tau = o(\frac{\log n}{\log \log n})$, this is $n^{o(1)}$. The preprocessing time is also $n^{o(1)}$.

It is a matter of simple arithmetic to extend a given fragment v of T, $|v| \leq \tau$, to a block T_i. We use the precomputed answers stored for $\mathrm{oid}(T_i)$ to determine the minimal suffix of v. We only need to translate the indices within T_i to indices within T before we return the answer. The following theorem summarizes our contribution for short fragments:

Theorem 4.8. *For every text T of length n and every parameter $\tau = o(\frac{\log n}{\log \log n})$ there exists a data structure of size $\mathcal{O}(\frac{n \log \tau}{\log n})$ which can answer in $\mathcal{O}(1)$ time* Minimal Suffix Queries *for fragments of length not exceeding τ. Moreover, it can be constructed in $\mathcal{O}(n)$ time.*

As noted at the beginning, this can be used to speed up queries for arbitrary fragments:

Theorem 4.9. *For every text T of length n there exists a data structure of size $\mathcal{O}(n)$ which can be constructed in $\mathcal{O}(n)$ time and answers* Minimal Suffix Queries *in $\mathcal{O}(1)$ time.*

5 Answering Generalized Minimal Suffix Queries

In this section we develop our data structure for Generalized Minimal Suffix Queries. We start with preliminary definitions and then we describe the counterparts of the three data structures presented in Section 4. Their query times are $\mathcal{O}(k^2 \log |v|)$, $\mathcal{O}(k^2 \log^* |v|)$, and $\mathcal{O}(k^2)$, respectively, i.e., there is an $\mathcal{O}(k^2)$ overhead compared to Minimal Suffix Queries.

We define a *k-fragment* of a text T as a concatenation $T[\ell_1..r_1] \cdots T[\ell_k..r_k]$ of k fragments of the text T. Observe that a k-fragment can be stored in $\mathcal{O}(k)$ space as a sequence of pairs (ℓ_i, r_i). If a string w admits such a decomposition using k' ($k' \leq k$) substrings, we call it a *k-substring* of T. Every k'-fragment (with $k' \leq k$) whose value is equal to w is called an *occurrence* of w as a k-substring of T. Observe that a substring of a k-substring w of T is itself a k-substring of T. Moreover, given an occurrence of w, one can canonically assign each fragment of w to a k'-fragment of T ($k' \leq k$). This can be implemented in $\mathcal{O}(k)$ time and referring to $w[\ell..r]$ in our algorithms, we assume that such an operation is performed.

Basic queries regarding k-fragments easily reduce to their counterparts for 1-fragments:

Observation 5.1. *The enhanced suffix array can answer queries (a), (b), and (c), in $\mathcal{O}(k)$ time if x and y are k-fragments of T.*

Generalized Minimal Suffix Queries can be reduced to the following auxiliary queries:

> Auxiliary Minimal Suffix Queries
> Given a fragment v of T and a k-fragment w of T, compute $\mathrm{MinSuf}(v, w)$ (represented as a $(k+1)$-fragment of T).

Lemma 5.2. *For every text T, the minimal suffix of a k-fragment v can be determined by k* Auxiliary Minimal Suffix Queries *(with $k' < k$) and additional $\mathcal{O}(k^2)$-time processing using the enhanced suffix array of T.*

Proof. Let $v = v_1 \cdots v_k$. By Observation 3.2(b), $\mathrm{MinSuf}(v) = \mathrm{MinSuf}(v_k)$ or for some i, $1 \leq i < k$, we have $\mathrm{MinSuf}(v) = \mathrm{MinSuf}(v_i, v_{i+1} \cdots v_k)$. Hence, we apply Auxiliary Minimal Suffix Queries to determine $\mathrm{MinSuf}(v_i, v_{i+1} \cdots v_k)$ for each $1 \leq i < k$. Observation 3.2(c) lets reduce computing $\mathrm{MinSuf}(v_k)$ to another auxiliary query. Having obtained k candidates for $\mathrm{MinSuf}(v)$, we use the enhanced suffix array to return the smallest among them using $k-1$ comparisons, each performed in $\mathcal{O}(k)$ time; see Theorem 2.1 and Observation 5.1. □

Fact 5.3. Auxiliary Minimal Suffix Queries *can be answered in* $\mathcal{O}(k \log |v|)$ *time using the enhanced suffix array of* T.

Proof. We apply Lemma 4.2 to determine $\Lambda(v)$, and then we compute the smallest string among $\{sw : s \in \Lambda(v)\}$. These strings are $(k+1)$-fragments of T and thus a single comparison takes $\mathcal{O}(k)$ time using the enhanced suffix array. □

Corollary 5.4. Generalized Minimal Suffix Queries *can be answered in* $\mathcal{O}(k^2 \log |v|)$ *time using the enhanced suffix array of* T.

5.1 $\mathcal{O}(k \log^* |v|)$-time Auxiliary Minimal Suffix Queries

Our data structure closely follows its counterpart described in Section 4.1. We define distinguished fragments in the same manner and provide a recursive algorithm based on Fact 4.4. However, for each distinguished fragment instead of applying Lemma 4.5, we build the following much stronger data structure. Its implementation is provided in Section 5.1.2.

Lemma 5.5. *Let* v *be a fragment of* T. *There exists a data structure of size* $\mathcal{O}(\log^2 |v|)$ *which answers the following queries in* $\mathcal{O}(k)$ *time: determine* $\mathrm{MinSuf}(v, w)$ *for a given* k*-fragment* w *of* T. *The data structure can be constructed in* $\mathcal{O}(\log^2 |v|)$ *time; it assumes the access to the enhanced suffix array of* T.

If $f(|v|) \geq |v|$ ($|v| \leq 2^{16}$), we use Fact 5.3 to compute $\mathrm{MinSuf}(v, w)$ in $\mathcal{O}(k \log |v|) = \mathcal{O}(k)$ time. Otherwise, we apply Fact 4.4 to decompose $v = uv'v''$ so that v' is distinguished, $|u| \leq |v'|$, and $|v''| \leq f(|v|)$, where $f(x) = 2^{\lfloor \log \log x \rfloor^2}$. The characterization of Observation 3.2 and Lemma 3.9 again gives three candidates for $\mathrm{MinSuf}(v, w)$: $\mathrm{MaxSuf}^R(u, v')v''w$, $\mathrm{MinSuf}(v', v''w)$, and $\mathrm{MinSuf}(v'', w)$. We determine the first using Lemma 4.1, the second using Lemma 5.5, while the third is computed recursively. The application of Lemma 5.5 takes $\mathcal{O}(k+1)$ time, since $v''w$ is a $(k+1)$-fragment of T. We return the best of the three candidates using the enhanced suffix array to choose it in $\mathcal{O}(k)$ time. Since $f(f(x)) = o(\log x)$, the depth of the recursion is $\mathcal{O}(\log^* |v|)$. This concludes the proof of the following result:

Theorem 5.6. *For every text* T *of length* n *there exists a data structure of size* $\mathcal{O}(n)$ *which answers* Auxiliary Minimal Suffix Queries *in* $\mathcal{O}(k \log^* |v|)$ *time and* Generalized Minimal Suffix Queries *in* $\mathcal{O}(k^2 \log^* |v|)$ *time. The data structure can be constructed in* $\mathcal{O}(n)$ *time.*

5.1.1 Rank queries in a collection of fragments

The crucial tool we use in the proof of Lemma 5.5 is a data structure constructed for a collection A of $W^{\mathcal{O}(1)}$ fragments of T to support $\mathrm{rank}_A(w)$ queries for arbitrary k-fragments w of T. Since it heavily relies on the compressed trie of fragments in A, we start by recalling several related concepts.

A trie is a rooted tree whose nodes correspond to prefixes of strings in a given family of strings A. If ν is a node, the corresponding prefix v is called the *value* of the node. The node whose value is v is called the *locus* of v.

The parent-child relation in the trie is defined so that the root is the locus of ε, while the parent ν' of a node ν is the locus of the value of ν with the last character removed. This character is the *label* of the edge from ν' and ν. In general, if ν' is a ancestor of ν, then label of the path from ν' to ν is the concatenation of edge labels on the path.

A node is *branching* if it has at least two children and *terminal* if its value belongs to A. A *compressed trie* is obtained from the underlying trie by dissolving all nodes except the root, branching nodes, and terminal nodes. Note that this way we compress paths of vertices with single children, and thus the number of remaining nodes becomes bounded by $2|A|$. In general, we refer to all preserved nodes of the trie as *explicit* (since they are stored explicitly) and to the dissolved ones as *implicit*. Edges of a compressed trie correspond to paths in the underlying tree and thus their labels are strings in Σ^+. Typically, these labels are stored as references to fragments of the strings in A.

Before we proceed with ranking a k-fragment in a collection of fragments, let us prove that fusion trees make it relatively easy to rank a suffix in a collection of $W^{\mathcal{O}(1)}$ suffixes.

Fact 5.7. *Let A be a set of $W^{\mathcal{O}(1)}$ suffixes of T. There exists a data structure of size $\mathcal{O}(|A|)$, which answers the following queries in $\mathcal{O}(1)$ time: given a suffix v of T, find a suffix $u \in A$ maximizing $\operatorname{lcp}(u, v)$. The data structure can be constructed in $\mathcal{O}(|A|)$ time; it assumes the access to the enhanced suffix array of T.*

Proof. Let $A = \{T[\ell_1..], \ldots, T[\ell_m..]\}$. We build a fusion tree storing $\{ISA[\ell_i] : 1 \leq i \leq m\}$ and during a query for $v = T[\ell..]$, we determine the predecessor and the successor of $ISA[\ell]$. We use the SA table to translate these integers into indices ℓ_{i_p} and ℓ_{i_s}. Since the order of $ISA[\ell_i]$ coincides with the lexicographic order of suffixes $T[\ell_i..]$, the suffixes $T[\ell_{i_p}..]$ and $T[\ell_{i_s}..]$ are the predecessor $\operatorname{pred}_A(v)$ and the successor $\operatorname{succ}_A(v)$, respectively. These are the two candidates for $u \in A$ maximizing $\operatorname{lcp}(u, v)$. We perform two longest common prefix queries and return the candidate for which the obtained value is larger. □

Lemma 5.8. *Let A be a set of $W^{\mathcal{O}(1)}$ fragments of T. There exists a data structure of size $\mathcal{O}(|A|^2)$, which answers the following queries in $\mathcal{O}(k)$ time: given a k-fragment v of T, determine $\operatorname{rank}_A(v)$. The data structure can be constructed in $\mathcal{O}(|A|^2)$ time; it assumes the access to the enhanced suffix array of T.*

Proof. Let $A = \{T[\ell_1..r_1], \ldots, T[\ell_m..r_m]\}$ and let $\mathcal{T}$ be the compressed trie of fragments in A. Note that $\mathcal{T}$ can be easily constructed in $\mathcal{O}(m \log m)$ time using the enhanced suffix array. For each edge we store a fragment of T representing its label and for each terminal node its rank in A. Moreover, for each explicit node ν of $\mathcal{T}$ we store pointers to the first and last (in pre-order) terminal nodes in its subtree as well as the following two components: a fusion tree containing the children of ν indexed by the first character of the corresponding edge label, and a data structure of Fact 5.7 for $\{T[\ell_i + d_\nu..] : \ell_i \in L_\nu\}$, where d_ν is the (weighted) depth of ν and L_ν contains ℓ_i whenever the locus of $T[\ell_i..r_i]$ is in the subtree of ν. Finally, for each ℓ_i we store a fusion tree containing (pointers to) all explicit nodes of $\mathcal{T}$ which represent prefixes of $T[\ell_i..]$, indexed by their (weighted) node depths. All these components can be constructed in $\mathcal{O}(m^2)$ time, with Theorem 2.2 applied to build fusion trees.

Let us proceed to the description of a query algorithm. Let $v = v_1 \cdots v_k$ be the decomposition of the given k-fragment into 1-fragments, and let $p_i = v_1 \cdots v_i$ for $0 \leq i \leq k$. We shall scan all v_i consecutively and after processing v_i, store a pointer to the (possibly implicit) node ν_i defined as the locus of the longest prefix of p_i present in $\mathcal{T}$. We start with $p_0 = \varepsilon$ whose locus is the root of $\mathcal{T}$. Therefore, it suffices to describe how to determine ν_i provided that we know ν_{i-1}.

If ν_{i-1} is at depth smaller than $|p_{i-1}|$, there is nothing to do, since $\nu_i = \nu_{i-1}$. Otherwise, we proceed as follows: Let ν be the nearest explicit descendant of ν_{i-1} ($\nu = \nu_{i-1}$ if ν_{i-1} is explicit), and let u be a fragment of T representing the label from ν_{i-1} to ν. First, we check if u is a proper prefix of v_i. If not, ν_i is on the same edge of $\mathcal{T}$ and its depth $|p_{i-1}| + |\operatorname{lcp}(u, v_i)|$. Thus, we may assume that u is a proper prefix of v_i. Let $v_i = uT[\ell..r]$. We make a query to the data structure of Fact 5.7 built for ν with $T[\ell..]$ as the query suffix. This lets us determine an index $\ell_j \in L_\nu$ such that $\operatorname{lcp}(T[\ell..], T[\ell_j + d_\nu..])$ is largest possible. This is also an index $\ell_j \in L_\nu$ which maximizes $D := \operatorname{lcp}(p_i, T[\ell_j..]) = d_\nu + \operatorname{lcp}(T[\ell..], T[\ell_j + d_\nu..])$. Consequently, ν_i represents a prefix of $T[\ell_j..]$ and the depth of ν_i does not exceed D. Thus, the nearest explicit ancestor of ν_i can be retrieved from the fusion tree built for ℓ_j as the node whose depth D' is the predecessor of D. If $D' < |p_i|$, we check if that explicit node has an outgoing edge whose label starts with $p_i[D' + 1] = T[\ell + D' + 1 - d_\nu]$. If not, ν_i is equal to the explicit node. Otherwise, ν_i is an implicit node on the found edge and its depth can be determined using a single longest common prefix query.

After processing the whole k-fragment v we are left with ν_k which is the locus of the longest prefix p of v present in $\mathcal{T}$. First, suppose that $p \neq v$ and let $c = v[|p|+1]$. Note that by definition of ν_k, this node does not have an outgoing edge labeled with c. If ν_k has no outgoing edge labeled with a character smaller then c, then the first terminal node of the subtree rooted at the leftmost child of ν_k represents the successor of v in A. We return its rank as the rank of v. Otherwise, we determine the edge going from ν_k to some node ν so that the edge label is smaller than c and largest possible. If ν_k is explicit, we use the fusion tree to determine ν. We observe that the predecessor of v in A is the rightmost terminal node in the subtree of ν and thus we return the rank stored at that node plus one. Thus, it remains to consider the case when $p = v$. In this case the leftmost terminal node in the subtree of ν_k is the successor of v in A, and thus we return the rank of that node. □

5.1.2 Proof of Lemma 5.5

Having developed the key component, we are ready to generalize Lemma 4.5.

Lemma 5.5. *Let v be a fragment of T. There exists a data structure of size $\mathcal{O}(\log^2 |v|)$ which answers the following queries in $\mathcal{O}(k)$ time: determine $\mathrm{MinSuf}(v, w)$ for a given k-fragment w of T. The data structure can be constructed in $\mathcal{O}(\log^2 |v|)$ time; it assumes the access to the enhanced suffix array of T.*

Proof. We use Lemma 4.2 to compute $\Lambda(v)$ in $\mathcal{O}(\log |v|)$ time. By Lemma 3.7, in order to find $\mathrm{MinSuf}(v, w)$, it suffices determine $\mathrm{rank}_{X(v)}(w)$. Moreover, by Lemma 3.7 and Observation 3.6, $\mathrm{rank}_{X(v)}(w)$ is equal to $\mathrm{rank}_{X'(v)}(w)$ or $\mathrm{rank}_{X'(v)}(w) - 1$, where $X'(v) = \{x_\lambda^{p_\lambda}, \ldots, x_m^{p_m}\}$ can be determined in $\mathcal{O}(\log |v|)$ time from $\Lambda(v)$. We build the data structure of Lemma 5.8 for $A = X'(v)$ so that we can determine $\mathrm{rank}_{X'(v)}(w)$ in $\mathcal{O}(k)$ time. This leaves two possibilities for $\mathrm{rank}_{X(v)}(w)$, i.e., for $\mathrm{MinSuf}(v, w)$. We simply need to compare $s_i w$, $s_{i+1} w$ for these two candidates suffixes $s_i, s_{i+1} \in \Lambda(v)$. Using the enhanced suffix array, this takes $\mathcal{O}(k)$ time. Consequently, the query algorithm takes $\mathcal{O}(k)$ time in total. In the preprocessing we need to compute $\Lambda(v)$ and the data structure of Lemma 5.8 for $A = X'(v)$, which takes $\mathcal{O}(\log |v| + |\Lambda(v)|^2) = \mathcal{O}(\log^2 |v|)$ time. The space consumption is also $\mathcal{O}(\log^2 |v|)$. □

5.2 $\mathcal{O}(k)$-time Auxiliary Minimal Suffix Queries

Like in Section 4.2, in order to improve the query time in the data structure of Theorem 5.6, we simply add a component responsible for computing $\mathrm{MinSuf}(v, w)$ for $|v| \leq \tau$ where $\tau = f(f(n)) = o(\frac{\log n}{\log \log n})$.

Again, we partition T into $\frac{n}{\tau}$ overlapping blocks T_i of length $m = \mathcal{O}(\tau)$, so that the number of blocks is much larger than the number of order-isomorphism classes of strings of length $\leq m$. Next, we precompute some data for each equivalence class and we reduce a query in T to a query in one of the blocks T_i.

While this approach was easy to apply for computing $\mathrm{MinSuf}(v)$ for a fragment v (with $|v| \leq \tau$), it is much more difficult for $\mathrm{MinSuf}(v, w)$ for a fragment v ($|v| \leq \tau$) and a k-fragment w. That is because w might be composed of fragments w_j starting in different blocks. As a workaround, we shall replace w by a similar (in a certain sense) k'-fragment of $T_i\$$ ($k' \leq k+1$) where T_i is a block containing v.

For $0 \leq i < \frac{n}{\tau}$, we define $T_i := T[i\tau + 1 .. \min(n, (i+3)\tau)]$. We determine $\mathrm{oid}(T_i\$)$ for each block using Fact 4.7. For each valid identifier we build the enhanced suffix array and for all fragment v we construct the set $\Lambda(v)$ along with the data structure of Lemma 5.5. In total, this data takes $\mathcal{O}(2^{\mathcal{O}(m \log m)} m^{\mathcal{O}(1)}) = n^{o(1)}$ space and time to construct.

Now, suppose that we are to compute $\mathrm{MinSuf}(v, w)$ where $|v| \leq \tau$ and w is a k-fragment of T. We determine the last block T_i containing v. Next, we shall try to represent w as a k-fragment of T_i. We will either succeed, or obtain a k'-fragment w' of T_i ($k' \leq k$) and a character $c \in \Sigma$ such that $w'c$ is a prefix of w but not a substring of v. In this case Lemma 3.10 states that $\mathrm{MinSuf}(v, w'c)$ suffices to determine two candidates for $\mathrm{MinSuf}(v, w)$.

We decompose $w = w_1 \cdots w_k$ into fragments and process them iteratively. Given a fragment w_j we shall either find an equal fragment of T_i or determine a fragment w'_j of T_i and a character $c \in \Sigma$ such that $w'_j c$ is a prefix of w_j but not a substring of v. Clearly, if we proceed to w_{j+1} in the first case and terminate in the second, at the end we successfully represent w or we find a k'-fragment $w' = w_1 \ldots w_{k'-1} w'_{k'}$ satisfying the

desired condition. Note that since v is a substring of $T[i\tau+1..\min(n,(i+2)\tau)]$, any substring of v, must occur in T at one of the positions in $R_i = \{i\tau+1,\ldots,\min(n,(i+2)\tau)\}$. Hence, for each block we build a data structure of Fact 5.7 for suffixes starting in R_i. Given w_j this lets us determine a position $\ell \in R_i$ such that $d_j = \mathrm{lcp}(T[\ell..], w_j)$ is largest possible. If $d_j = |w_j|$ and $d_j \le \tau$, we have found w_j occurring as a substring of T_i. Otherwise, we set $w'_j = w[1..\min(d_j,\tau)]$, which is a substring of T_i, and $c = w[|w'_j|+1]$. Clearly, $w'_j c$ is a prefix of w_j, so we shall only prove that it is not a substring of v. If $d_j \ge \tau$, then simply $|w'_j c| > \tau \ge |v|$. Otherwise, by the choice of ℓ maximizing $d_j = \mathrm{lcp}(T[\ell..], w_j)$ among $\ell \in R_i$, the string $w'_j c$ cannot occur at any position in R_i and in particular it cannot be a substring of v.

If the described procedure succeeds in finding a k-fragment of T_i equal to w, we simply apply the data structure of Lemma 5.5 built for v to determine $\mathrm{MinSuf}(v,w)$ in $\mathcal{O}(k)$ time. Thus, we may assume that this is not the case and it returns a k'-fragment w' and a character c. As already mentioned, having computed $\mathrm{MinSuf}(v,w'c)$, we can determine $\mathrm{MinSuf}(v,w)$ just by comparing the two candidates with the enhanced suffix array. If c occurs in T_i, then $w'c$ is a $(k'+1)$-fragment of T_i and we may use Lemma 5.5 to compute $\mathrm{MinSuf}(v,w'c)$. Otherwise, we replace c by its successor among letters occurring in $T_i\$$. The successor can be computed in constant time provided that for each block we store a fusion tree of all characters occurring in $T_i\$$ (mapping each character to a sample position). To see that replacing c by its successor c' does not change the answer, it is enough to note that Lemma 3.7 expresses $\mathrm{MinSuf}(v,w'c)$ in terms of $\mathrm{rank}_{X(v)}(w'c)$, where $X(v)$ consists of infinite strings composed of characters of v (which are automatically present in T_i).

Theorem 5.9. *For every text T of length n and every parameter $\tau = o(\frac{\log n}{\log\log n})$ there exists a data structure of size $\mathcal{O}(n)$ which answers* Auxiliary Minimal Suffix Queries *in $\mathcal{O}(k)$ time if $|v| \le \tau$. The data structure can be constructed in $\mathcal{O}(n)$ time.*

This was the last missing ingredient needed to obtain the main result of this paper.

Theorem 5.10. *For every text T of length n there exists a data structure of size $\mathcal{O}(n)$ which answers* Auxiliary Minimal Suffix Queries *in $\mathcal{O}(k)$ time and* Generalized Minimal Suffix Queries *in $\mathcal{O}(k^2)$ time. The data structure can be constructed in $\mathcal{O}(n)$ time.*

6 Applications

As already noted in [5], Minimal Suffix Queries queries can be used to compute Lyndon factorization. For fragments of T, and in general $k = \mathcal{O}(1)$, we obtain an optimal solution:

Corollary 6.1. *For every text T of length n there exists a data structure of size $\mathcal{O}(n)$ which given a k-fragment v of T determines the Lyndon factorization $v = v_1^{q_1}\ldots v_m^{q_m}$ in $\mathcal{O}(k^2 m)$ time. The data structure takes $\mathcal{O}(n)$ time to construct.*

Our main motivation of introducing Generalized Minimal Suffix Queries, however, was to answer Minimal Rotation Queries, for which we obtain constant query time after linear-time preprocessing. This is achieved using the following observation; see [11]:

Observation 6.2. *The minimal cyclic rotation of v is the prefix of* $\mathrm{MinSuf}(v,v)$ *of length $|v|$.*

Theorem 6.3. *For every text T of length n there exists a data structure of size $\mathcal{O}(n)$ which given a k-fragment v of T determines the lexicographically smallest cyclic rotation of v in $\mathcal{O}(k^2)$ time. The data structure takes $\mathcal{O}(n)$ time to construct.*

Using Minimal Rotation Queries, we can compute the Karp-Rabin fingerprint [23] of the minimal rotations of a given fragment v of T (or in general, of a k-fragment). This can be interpreted as a computing fingerprints up to cyclic equivalence, i.e., evaluating a function h such that $h(\ell,r) = h(\ell',r')$ if and only if $T[\ell..r]$ and $T[\ell'..r']$ are cyclically equivalent.

Consequently, we are able, for example, to count distinct substrings of T with a given exponent $1+1/\alpha$. They occur within runs or α-gapped repeats, which can be generated in time $\mathcal{O}(n\alpha)$ [7, 13, 18] and classified

using Minimal Rotation Queries according to the cyclic equivalence class of their period. For a fixed equivalence class the set of substrings generated by a single repeat can be represented as a cyclic interval, and the cardinality of a union of intervals is simple to determine; see also [12], where this approach was used to count and list squares and, in general, substrings with a given exponent 2 or more.

Acknowledgements I would like to thank the remaining co-authors of [5], collaboration with whom on earlier results about minimal and maximal suffixes sparked some of my ideas used in this paper. Special acknowledgments to Paweł Gawrychowski and Tatiana Starikovskaya for numerous discussions on this subject.

References

[1] Amihood Amir, Alberto Apostolico, Gad M. Landau, Amit Levy, Moshe Lewenstein, and Ely Porat. Range LCP. *Journal of Computer and System Sciences*, 80(7):1245–1253, 2014.

[2] Amihood Amir, Moshe Lewenstein, and Sharma V. Thankachan. Range LCP queries revisited. In Costas S. Iliopoulos, Simon J. Puglisi, and Emine Yilmaz, editors, *String Processing and Information Retrieval, SPIRE 2015*, volume 9309 of *LNCS*, pages 350–361. Springer, 2015.

[3] Alberto Apostolico and Maxime Crochemore. Optimal canonization of all substrings of a string. *Information and Computation*, 95(1):76–95, 1991.

[4] Alberto Apostolico and Maxime Crochemore. Fast parallel Lyndon factorization with applications. *Theory of Computing Systems*, 28:89–108, 1995.

[5] Maxim Babenko, Paweł Gawrychowski, Tomasz Kociumaka, Ignat Kolesnichenko, and Tatiana Starikovskaya. Computing minimal and maximal suffixes of a substring. *Theoretical Computer Science*, 2015. In press, 10.1016/j.tcs.2015.08.023.

[6] Maxim Babenko, Paweł Gawrychowski, Tomasz Kociumaka, and Tatiana Starikovskaya. Wavelet trees meet suffix trees. In Piotr Indyk, editor, *26th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015*, pages 572–591. SIAM, 2015.

[7] Hideo Bannai, Tomohiro I, Shunsuke Inenaga, Yuto Nakashima, Masayuki Takeda, and Kazuya Tsuruta. The "runs" theorem. *CoRR*, abs/1406.0263v7, 2015.

[8] Michael A. Bender and Martin Farach-Colton. The LCA problem revisited. In Gaston H. Gonnet, Daniel Panario, and Alfredo Viola, editors, *Latin American Symposium on Theoretical Informatics, LATIN 2000*, volume 1776 of *LNCS*, pages 88–94. Springer Berlin Heidelberg, 2000.

[9] Kuo Tsai Chen, Ralph Hartzler Fox, and Roger Conant Lyndon. Free differential calculus, IV. The quotient groups of the lower central series. *The Annals of Mathematics*, 68(1):81–95, 1958.

[10] Graham Cormode and S. Muthukrishnan. Substring compression problems. In *16th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2005*, pages 321–330. SIAM, 2005.

[11] Maxime Crochemore, Christophe Hancart, and Thierry Lecroq. *Algorithms on Strings*. Cambridge University Press, New York, NY, USA, 2007.

[12] Maxime Crochemore, Costas S. Iliopoulos, Marcin Kubica, Jakub Radoszewski, Wojciech Rytter, and Tomasz Waleń. Extracting powers and periods in a word from its runs structure. *Theoretical Computer Science*, 521:29–41, 2014.

[13] Maxime Crochemore, Roman Kolpakov, and Gregory Kucherov. Optimal searching of gapped repeats in a word. *CoRR*, abs/1509.01221, 2015.

[14] Maxime Crochemore and Dominique Perrin. Two-way string-matching. *Journal of the ACM*, 38(3):650–674, July 1991.

[15] Jacqueline W. Daykin, Costas S. Iliopoulos, and William F. Smyth. Parallel RAM algorithms for factorizing words. *Theoretical Computer Science*, 127(1):53–67, 1994.

[16] Jean-Pierre Duval. Factorizing words over an ordered alphabet. *Journal of Algorithms*, 4(4):363–381, 1983.

[17] Michael L. Fredman and Dan E. Willard. Surpassing the information theoretic bound with fusion trees. *Journal of Computer and System Sciences*, 47(3):424–436, 1993.

[18] Paweł Gawrychowski, Tomohiro I, Shunsuke Inenaga, Dominik Köppl, and Florin Manea. Efficiently finding all maximal α-gapped repeats. *CoRR*, abs/1509.09237, 2015. Accepted to STACS 2016.

[19] Torben Hagerup. Sorting and searching on the word RAM. In Michel Morvan, Christoph Meinel, and Daniel Krob, editors, *Symposium on Theoretical Aspects of Computer Science, STACS 1998*, volume 1373 of *LNCS*, pages 366–398. Springer, Berlin Heidelberg, 1998.

[20] Dov Harel and Robert Endre Tarjan. Fast algorithms for finding nearest common ancestors. *SIAM Journal on Computing*, 13(2):338–355, 1984.

[21] Tomohiro I, Yuto Nakashima, Shunsuke Inenaga, Hideo Bannai, and Masayuki Takeda. Efficient lyndon factorization of grammar compressed text. In Johannes Fischer and Peter Sanders, editors, *Combinatorial Pattern Matching, CPM 2013*, volume 7922 of *LNCS*, pages 153–164. Springer, 2013.

[22] Tomohiro I, Yuto Nakashima, Shunsuke Inenaga, Hideo Bannai, and Masayuki Takeda. Faster Lyndon factorization algorithms for SLP and LZ78 compressed text. In Oren Kurland, Moshe Lewenstein, and Ely Porat, editors, *String Processing and Information Retrieval, SPIRE 2013*, volume 8214 of *LNCS*, pages 174–185. Springer International Publishing Switzerland, 2013.

[23] Richard M. Karp and Michael O. Rabin. Efficient randomized pattern-matching algorithms. *IBM Journal of Research and Development*, 31(2):249–260, 1987.

[24] Orgad Keller, Tsvi Kopelowitz, Shir Landau Feibish, and Moshe Lewenstein. Generalized substring compression. *Theoretical Computer Science*, 525:45–54, 2014.

[25] Jinil Kim, Peter Eades, Rudolf Fleischer, Seok-Hee Hong, Costas S. Iliopoulos, Kunsoo Park, Simon J. Puglisi, and Takeshi Tokuyama. Order-preserving matching. *Theoretical Computer Science*, 525:68–79, 2014.

[26] Tomasz Kociumaka, Jakub Radoszewski, Wojciech Rytter, and Tomasz Waleń. Internal pattern matching queries in a text and applications. In Piotr Indyk, editor, *26th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015*, pages 532–551. SIAM, 2015.

[27] Marcin Kubica, Tomasz Kulczyński, Jakub Radoszewski, Wojciech Rytter, and Tomasz Waleń. A linear time algorithm for consecutive permutation pattern matching. *Information Processing Letters*, 113(12):430–433, 2013.

[28] Roger Conant Lyndon. On Burnside's problem. *Transactions of the American Mathematical Society*, 77(2):202–215, 1954.

[29] Udi Manber and Eugene W. Myers. Suffix arrays: A new method for on-line string searches. *SIAM Journal on Computing*, 22(5):935–948, 1993.

[30] Marcin Mucha. Lyndon words and short superstrings. In Sanjeev Khanna, editor, *24th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013*, pages 958–972. SIAM, 2013.

[31] Manish Patil, Rahul Shah, and Sharma V. Thankachan. Faster range LCP queries. In Oren Kurland, Moshe Lewenstein, and Ely Porat, editors, *String Processing and Information Retrieval, SPIRE 2013*, volume 8214 of *LNCS*, pages 263–270. Springer International Publishing, 2013.

[32] Mihai Pătraşcu and Mikkel Thorup. Dynamic integer sets with optimal rank, select, and predecessor search. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014*, pages 166–175. IEEE Computer Society, 2014.